CUSTOMER NO. 23932

PATENT APPLICATION
Docket #61170-27USPX

EXPRESS MAIL Mailing Label No: EV 227 870 465 US

Date of Deposit 01-26-04

Type or Print Name MARGO BARBARASH

Signature

PROCESS FOR REFRESHING A DYNAMIC RANDOM ACCESS MEMORY AND CORRESPONDING DEVICE

PRIORITY CLAIM

The present application claims priority from French Application for Patent No. 03 01005 filed January 29, 2003, the disclosure of which is hereby incorporated by reference.

BACKGROUND OF THE INVENTION

Technical Field of the Invention

[1] The present invention relates to dynamic random access memories, that is to say those requiring periodic refreshing of the data contained in the memory cells of these memories. Also, the invention relates more particularly to the process of refreshing these dynamic random access memories.

[2] The invention applies advantageously, but without limitation, to cellular mobile telephones which incorporate dynamic random access memories.

Description of Related Art

[3] Third-generation cellular mobile telephones will require the integration of large quantities of memory. However, the cost of the product must remain low. The use of dynamic random access memories (DRAMs), in place of the static random access memories (SRAMs) as are currently used, allows for this rise in memory capacity to be achieved at a low cost.

[4] However, an important constraint in this type of application is the low electrical consumption while the telephone is on standby, so as not to discharge the batteries too quickly. However, although dynamic random access memories have a smaller static leakage current than static random access memories, they need to be refreshed continuously if the data has to be preserved in standby mode. This refreshing requires energy consumption that it is important to minimize.

[5] The refresh frequency is given by the number of memory pages to be refreshed and by the retention time of the memory. This retention time is essentially related to the junction leakages of the transistors of the memory cells. Now, these leakages vary essentially as a function of temperature.

[6] At present, dynamic random access memories are refreshed at the frequency corresponding to the worst operating case, that is to say to the maximum temperature of the application.

[7] The article by Jae-Yoon Sim et al., entitled "Double Boosting Pump, Hybrid Current Sense Amplifier, and Binary Weighted Temperature Sensor Adjustment Schemes for 1.8

V 128 Mb Mobile DRAMs," 2002 Symposium on VLSI Circuits Digest of Technical Papers, describes a temperature measurement system integrated onto the chip supporting the dynamic random access memory, and which acts on the refresh frequency. However, this system requires a tailoring of the temperature sensor at the on-board test level, this possibly complicating this test and therefore making it more expensive. Furthermore, a little current is consumed by the temperature sensor itself.

[8] There is a need to optimize the refresh frequency of a dynamic random access memory, and to do so in a manner which is particularly simple to implement and without using an external sensor.

SUMMARY OF THE INVENTION

[9] The present invention proposes a process for refreshing a dynamic random access memory in which the retention time of all the memory cells of the memory is measured continuously and dynamically, and the result of this measurement is used to regulate the refresh period of the memory.

[10] Use of all the memory cells of the memory makes it possible to adjust the refresh period of the memory by using the actual retention time of the entire memory, and not by using statistical characterizations of the ratio of mean retention time of the cells to minimum retention time over the memory, and also without the use of any other external means, such as a temperature sensor.

[11] According to one mode of implementation of the invention, the continuous and dynamic measurement of the retention time of all the cells comprises:

successive selections of different groups of cells of the memory, the so-called test cells, in such a way as to scan the entire set of cells of the memory,

successive elementary measurements of the retention times of each group of selected test cells, and

successive refreshings of the unselected cells.

[12] Furthermore, the said successive elementary measurements are performed at a lower measurement frequency than the refresh frequency of the other cells of the memory, in such a way as to "refresh" the groups of test cells more slowly than the other cells of the memory.

[13] When the memory is organized by pages, each group of test cells is preferably an integer number of pages, for example one page or several pages.

[14] A particularly simple way of performing the phase of elementary measurement of the retention time of a group of selected test cells may comprise:

a writing to these selected test cells of a predetermined test content corresponding to the charging of all the capacitors of the test cells, and

then a reading of the test cells and a metering of the number of errors in reading the content of the test cells having regard to the said test content.

[15] Moreover, the measurement of the retention time of all the cells comprises for example an at least partial accumulation of the successive numbers of errors.

[16] Finally, between the writing of the test content to the test cells and the reading of these test cells, at least two refreshings of the other cells of the memory are performed.

[17] In general, the regulating of the refresh period is performed at the end of the test of the entire set of cells of the memory. However, it is also possible, and this is the reason why one speaks of at least partial accumulation, for the regulating of the refresh period to be able to be performed as soon as a certain number of errors has been detected, without waiting for the end of the test of the entire memory, in particular if one wishes to be able to react faster to a temperature variation.

[18] When pages of the memory do not contain any useful data and are used as test pages, there is no reason to back up the content thereof. On the other hand, when memory pages do contain useful data and are in their turn selected as test cells, their content must be backed up.

[19] This is why the elementary measurement phase furthermore comprises, for certain at least of the groups of selected test cells, a backup of the content of these test cells, and a restoring of the content of the test cells with the backed up content once the metering of the number of errors has been performed.

[20] The backup can be performed in a predetermined part of the memory, for example in pages containing no useful data, or else if one does not wish to place any constraint on the use of certain pages of the memory, in an external backup memory. Of course, in this case, this external memory will have to contain at least as many pages as are tested in a retention time measurement cycle.

[21] When two buffer memories connected to the dynamic memory are provided, the elementary measurement phase advantageously comprises a writing of the test content into a first buffer memory before writing to the test cells. Furthermore, the reading of the content of the test

cells comprises a writing of the content of these test cells into the second buffer memory and then a reading of the second buffer memory.

[22] This allows a further decrease in the consumption of current. Specifically, the periodic reloading of the test content is then done as a single operation of the buffer memory in the test page, instead of having to systematically rewrite all the words of the memory page.

[23] Moreover, when just a single memory page is used as test page, it is advantageously possible to use one of the buffer memories as external backup memory.

[24] Thus, according to one mode of implementation of the invention, the content of a few pages of the dynamic random access memory is backed up, then one tries to refresh them less quickly, for example two times less quickly, and one observes whether this does or does not cause errors. The operation is repeated on the entire memory. Depending on the number of errors that have appeared on the pages refreshed less often, the refresh period is decreased or increased. Thus, the memory self-adjusts its refresh period to what is necessary for it.

[25] The number of test pages is for example chosen in such a way that the time required to test the retention of all the pages of the memory is sufficiently small as compared with the times of changes of temperature of the system. Thus, by way of indication, if one considers a memory comprising 4096 pages with minimum retention equal to 32 ms, and if one wishes to gain, at typical temperature, by a factor of 4 with regard to consumption due to refreshing at the minimum temperature, hence to have a maximum refresh period of 4 x 32 ms, and if one also assumes that the memory is refreshed twice as often as the test pages, then 256 ms (2 x 4 x 32) is necessary per set of test pages. So that the refresh period is evaluated every minute, it is consequently necessary to use around 16 test pages (4096/(60 s/256 ms)).

[26] As far as the regulating of the refresh period is concerned, several solutions are possible.

[27] According to a first possibility, the regulating of the refresh period of the memory may comprise

a comparison of the number of accumulated errors with a low threshold and a high threshold,

an increasing of the refresh period if the number of errors is less than the low threshold,

a decreasing of the refresh period if the number of errors is greater than the high threshold, and

a non-modification of the refresh period if the number of errors is greater than or equal to the low threshold and less than or equal to the high threshold.

[28] Another solution, which is simpler, for the regulating of the refresh period of the memory can comprise comparing the number of accumulated errors with a single threshold, and in increasing the refresh period if the number of errors is less than or equal to the single threshold and in decreasing the refresh period if the number of errors is greater than or equal to the single threshold.

[29] It is moreover preferable to make provision for a minimum limit value and, more particularly, a maximum limit value for the refresh period. Specifically, it is better to test the retention of the cells at a rate which remains fast compared with that of the movements in the temperature.

[30] Although the process according to the invention can be applied at any moment, it is more particularly intended to be applied during a standby mode. Thus, when the memory is

incorporated into an apparatus possessing a standby mode and an active mode of operation, such as for example a cellular mobile telephone, the mean retention time of the test cells is advantageously measured and the refresh period is regulated, at least in the course of the standby mode.

[31] The invention also makes it possible in particular to reduce the refresh frequency, hence the consumption in standby mode, even at the maximum operating temperature.

[32] The principle of such a mode of implementation is to tag the pages having the least retention, to note them, and to refresh them more often than the others. Specifically, generally, more than 90% of the pages of the memory may be refreshed two to four times less often than is necessary for the cells having a minimum retention.

[33] Stated otherwise, the measurement of the retention time of all the cells of the memory being performed cyclically, there is provision, according to one mode of realization of the invention, that in the course of a measurement cycle, for example the first, the cells of the memory having a lower retention are tagged. Then, in the course of the next cycle or cycles, these so-called low cells are refreshed more often than the other cells of the memory.

[34] The tagging of the "low" cells can also be performed regularly in the course of a current cycle of regulation of the memory refresh period.

[35] The invention also proposes a dynamic random access memory device, comprising a dynamic random access memory and means for refreshing the memory.

[36] According to a general characteristic of the invention, the device furthermore comprises auxiliary processing means able to continuously and dynamically measure the

retention time of all the memory cells of the memory, and to use the result of this measurement to regulate the refresh period of the memory.

[37] According to one embodiment of the invention, the auxiliary processing means comprise

selection means able to perform successive selections of different groups of cells of the memory, the so called test cells, in such a way as to scan the entire set of cells of the memory, and

measurement means able to perform successive elementary measurements of the retention times of each group of selected test cells.

[38] Moreover, the refresh means are able to perform successive refreshings of the unselected cells, and the auxiliary processing means comprise control means able to activate the measurement means less often than the refresh means, in such a way as to "refresh" the groups of test cells more slowly than the other cells of the memory.

[39] According to one embodiment of the invention, the measurement means comprise

storage means able to store a predetermined test content corresponding to the charging of all the capacitors of the test cells,

writing means able to write the test content into the selected test cells,

reading means able to read the content of the test cells,

metering means able to meter the number of errors in reading the content of the test cells having regard to the said test content,

accumulation means able to perform an at least partial accumulation of the successive numbers of errors.

[40] And between the writing of the test content to the test cells and the reading of these test cells, the control means are able to activate the refresh means at least twice.

[41] According to one embodiment of the invention, the measurement means furthermore comprise backup means able to perform a backup of the content of the test cells, and restoration means able to perform a restoration of the content of the test cells with the backed up content once the metering of the number of errors has been performed. Also the control means are able to activate the backup and restoration means for certain at least of the groups of selected test cells.

[42] According to one embodiment of the invention, the auxiliary processing means are able to cyclically perform the measurement of the retention time of all the cells of the memory. In the course of a measurement cycle, the auxiliary processing means are able to tag the cells of the memory having a lower retention, and in the course of the next cycle or cycles the refresh means are able to refresh these so-called low cells more often than the other cells of the memory.

[43] The device according to the invention is advantageously embodied in the form of an integrated circuit.

[44] The invention is also aimed at an apparatus possessing a standby mode and an active mode of operation, and incorporating a device as defined hereinabove, the auxiliary processing means being able to measure the retention time of all the cells of the memory and regulate the refresh period, at least in the course of the standby mode.

[45] This apparatus may be a component of a wireless communication system, for example a cellular mobile telephone.

BRIEF DESCRIPTION OF THE DRAWINGS

[46] A more complete understanding of the method and apparatus of the present invention may be acquired by reference to the following Detailed Description when taken in conjunction with the accompanying Drawings wherein:

[47] FIGURE 1 diagrammatically illustrates a random access memory according to the invention within which the test cells are grouped together in rows;

[48] FIGURE 2 illustrates in greater detail but still diagrammatically a memory device according to the invention and more particularly the auxiliary processing means associated with the random access memory according to the invention;

[49] FIGURE 3 represents a diagrammatic flowchart of a mode of implementation of the process according to the invention; and

[50] FIGURE 4 illustrates another mode of implementation of the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

[51] In FIGURE 1, the reference MMV designates a dynamic random access memory according to the invention whose memory plane PM comprises a matrix array of memory cells CL typically organized in rows RW and columns CLN. Each memory cell generally comprises a transistor and a capacitor. Additionally, in a conventional manner known per se, the memory plane is connected to a row decoder DCDL and to a column decoder (which are not represented here for simplifying purposes). Finally, in the example described here, the memory MMV comprises two buffer memories CH0 and CH1 connected to the memory plane PM.

[52] The architecture of such a memory equipped with two buffer memories is, for example, described in European patent application No. 952 587 (which is incorporated by reference), and will make possible, as will be seen in greater detail hereinbelow, a preferred implementation of the invention.

[53] In a general manner, according to the invention, the retention time of all the cells of the memory will be measured continuously and dynamically on the chip (integrated circuit) containing the memory MMV, and the refresh period of this memory will be adjusted accordingly.

[54] In the example described herein, the memory is organized into memory pages, a page corresponding to a line of words.

[55] Also, before returning in greater detail to the algorithm for adjusting the refresh period, we shall forthwith describe the basic outline of a mode of implementation thereof.

[56] More precisely, the content of one or more pages of this memory, that are then dubbed test page(s), is backed up and then one tries to refresh them less quickly, for example two times less quickly than the remaining pages. One observes whether this does or does not cause errors. The operation is repeated on the entire memory, changing test page each time. Depending on the number of errors that have appeared on the pages refreshed less often, the refresh period is decreased or increased. Thus, the memory self-adjusts its refresh period to what is necessary for it.

[57] In order to implement the mechanism for regulating the refresh frequency of the random access memory MMV, the memory device DMV according to the invention comprises (FIGURE 2), in addition to the memory MMV, auxiliary processing means MAT whose

structure and functionalities will now be described in greater detail. In a particular application of the invention, this device is incorporated into a cellular mobile telephone TP.

[58] The auxiliary processing means MAT comprise storage means RGT here formed of a register, and containing a predetermined test content PT1 corresponding to the charging of all the capacitors of the current test cells. More precisely, this test content PT1 comprises logic values that are chosen in such a way that, when they are written to the memory cells, the corresponding capacitors of these memory cells are charged to the voltage that corresponds to the worst case of retention. In the present case this voltage is equal to the supply voltage. That said, in other types of memory, for example based on PMOS transistors, this voltage may be the ground. Additionally, the values of these test bits depend on the manner in which the memory is embodied (even column, odd column, etc).

[59] The test content PT1 may be delivered on the input data bus DI of the memory by way of a multiplexer MUX1 controlled by a control signal. The other input of the multiplexer MUX1 receives the useful data to be written to the memory and emanating from a conventional controller CTLN dedicated to the normal mode of operation of the memory. This controller CTLN is moreover connected to the output data bus DO of the memory. Finally, it generates a refresh (RF) signal comprising the refresh commands as well as the memory control and address signals. Thus, all the cells of the memory are periodically refreshed, the period between two page refresh cycles being designated here by Tref.

[60] Comparison means CMP1 receive on the one hand the content of the register RGT, that is to say the test content PT1 and, on the other hand, the content of the test cells, by way of the output bus DO. The comparison means CMP1 will compare, at predetermined

instants, the content of the current test cells with the test content PT1. These comparison means CMP1 will, in combination with accumulation means ACC, make it possible to meter the number of errors in reading the current test cells, and also to accumulate the number of errors of the successive test pages during the complete scanning of the memory MMV. This number of errors NBE is stored in a register RG2.

[61] The structure used in this embodiment to meter the number of errors NBE is compatible with a reading of a word line in blocks of bits. More precisely, for example, when a word line contains 1024 bits, the 1024 bits are read in blocks of 64 bits, this requiring 16 accesses to the memory.

[62] Several modes of realization are then possible for performing the regulation of the refresh period by knowing the number NBE. One of these modes of realization is that illustrated in the remainder of FIGURE 2 and that will now be described.

[63] The number of errors NBE is compared, in comparison means CMP2, with a high threshold SH1 and with a low threshold SB1. The values of the high and low thresholds, SH1 and SB1, will be determined by the person skilled in the art as a function for example of the envisaged application and of the type of memory used.

[64] The output of the comparison means CMP2 is linked to the input of decision logic MRG which acts as means for regulating the refresh period Tref. More precisely, in a general manner, if the number of errors NBE is less than the low threshold SB1, the refresh period Tref is increased. If this number of errors is greater than the high threshold SH1, this period is decreased. If it is less than or equal to the high threshold and greater than or equal to the low threshold, the refresh period Tref is left unchanged.

[65] Hardware-wise, this can be achieved by a multiplexer MUX2 controlled by the regulating signal SRG emanating from the means MRG. This multiplexer MUX2 possesses three inputs. The middle input is linked to the output of the register RGT2 which contains the value of the refresh period Tref. Another input of the multiplexer is linked to the output of an adder. Such an adder possesses a first input linked to the output of the register RGT2 and a second input linked to a constant CH. The multiplexer MUX2 also possesses a third input linked to the output of a subtractor. This subtractor possesses a first input linked to the output of the register RGT2 and a second input linked to a constant CB. The output of the multiplexer MUX2 is linked to the input of the register RGT2.

[66] Thus, if the refresh period should remain unchanged, it is the middle input of the multiplexer MUX2 which is selected. On the other hand, if the refresh period should be increased, it is the left input (in FIGURE 2) which is chosen, the refresh period then being incremented by the constant CH. If the refresh period should be decreased, it is the right input of the multiplexer MUX2 which is selected, the refresh period then being decreased by the constant CB.

[67] This said, two other registers are also provided, respectively comprising a lower limit value SB2 and a limit value SH2 for the refresh period Tref. Also, the regulating of the refresh period just described is applied while the refresh period lies between these two limit values SB2 and SH2. On the other hand, if the refresh period is less than the limit value SB2, the regulating of the refresh period can only take the form of an increase in the refresh period. Likewise, if the refresh period is greater than the limit value SH2, then the regulating of the refresh period can only take the form of a decrease in the refresh period. In practice, on

initialization, a refresh period corresponding to the maximum refresh frequency having regard to the worst temperature case will be chosen. This initial value of the period then corresponds to the minimum limit value SB2. By way of indication, for a temperature of 85°C, the worst case retention time corresponds to 32 ms.

[68] As a variant, it is also possible to use a very simple refresh period modification decision algorithm using just a single threshold for the comparison of the number of errors NBE, for example a threshold equal to zero. More precisely, if there is no error in the entire memory, the value of the refresh period is increased (this new value having to be less than twice the previous value in the case where the test pages are "refreshed" two times less quickly than the remainder of the memory), otherwise it is decreased, with, as in the variant just described, a lower limit value SB2 and an upper limit value SH2.

[69] In this case, the logic for counting the number of errors NBE, rectangle LGC demarcated by dashes in FIGURE 2, reduces to a simple logical OR gate in regard to the error bits, and to a single flip-flop, set to 1 if at least one error is detected, and reset to zero by the state machine FSM at each main cycle of the algorithm (that is to say when the test pages again become those used at the start of the cycle).

[70] As just mentioned, in addition to the means just described, the auxiliary processing means MAT also comprise control means, in the form of a finite state machine FSM, and it is this machine FSM which will sequence the refreshing of the test cells and trigger the calculation of the number of errors NBE as well as the regulating of the refresh period Tref.

[71] To describe the various states of this machine FSM, reference will now be made more particularly to FIGURE 3. In this figure, and in this exemplary implementation, N

designates the memory page address, P designates the address of a word in a page and Q designates the number of refresh cycles between the writing of the test content into a test page and the reading of the content of this test page with a view to the metering of the errors. Q is for example equal to 2. Moreover, the chart of FIGURE 3 corresponds, for simplifying reasons, to the case where the number of test pages is equal to 1. Also, Ntest denotes the address of the current test page. Finally, T denotes the current time, counted as a number of cycles, since the last refresh.

[72] It is also assumed, in this mode of implementation, that the first page of the memory contains no useful data. It will consequently be seen that it is in this first page that the content of the subsequent pages will be backed up before they are tested.

[73] On entry to the standby mode, the first page of the memory is used as test page (Ntest equals zero; step 30). The variables N, P and q are also initialized to zero (step 31). As indicated hereinabove, the refresh period Tref is set to the minimum value, corresponding to the maximum temperature of use of the system. When P is equal to Pmax, that is to say when all the test words have been written to the buffer memory CH0, the content of the buffer memory CH0 is transferred (step 33) into the test page of the memory MMV.

[74] As Ntest is equal to zero, we go directly to step 35, in which N is incremented by one unit and then we proceed, in step 36, to the periodic refreshing of all the other pages of the random access memory. Thereafter, N is reinitialized to zero in step 37 and q is incremented.

[75] Since q is less than Q (here taken equal for example to 2), we return to step 35 and proceed to a second refreshing of the pages of the memory other than the test page. When q is equal to Q we then proceed to the "refreshing" of the test cells of the test page. In fact, here

the term "refreshing" is used improperly in respect of a test cell although it is not in reality a refreshing. More precisely, in step 38 the content of the test cells is read and stored in the buffer memory CH1 and then in step 39 the content of the buffer memory CH1 is read.

[76] We then proceed to the counting of the reading errors, by comparing the content thus read with the test content PT1.

[77] Thus, the person skilled in the art will have observed that the test cells are "refreshed" Q times less often than the other cells of the random access memory.

[78] By using two buffer memories the consumption of current can be further decreased since the test content is loaded into one of the buffer memories, the other buffer memory being used for reading. Also, the periodic reloading of the test content PT1 is then done in a single so-called "write back" operation of the buffer memory into the page, instead of having to systematically rewrite all the words of the page.

[79] Thereafter, one proceeds to a restoration of the content of the test page. In fact, this step 40 is not necessary in the present case for the first test page, but we shall see that it is necessary for the subsequent test pages.

[80] Next, Ntest is incremented by one unit (step 41) so that the page having the address 1 in the memory becomes in its turn the test page. One then proceeds to a backup (step 42) of the content of the new test page, for example here by transferring its content into the page with address zero. Next, after having reinitialized the variables N and q to zero (step 43), we return to step 33 in which the new test page is loaded with the test content PT1.

[81] For a test page having the address Ntest different from zero, steps 34, 35, 36 and 37 are performed in succession, during which we proceed to a refreshing of the pages of the

memory other than the test page. Then, since Q is greater than or equal to 1, one or more refreshings of the memory are again performed (as a function of the number Q) with the exception of the test page. Next, on completion of these refreshes, and after having performed the transfer 38 of the content of the test page into the buffer memory CH1, we again proceed to the counting of the errors related to the reading of this test page. Moreover, this new number of errors is accumulated with the previous number of errors corresponding to the previous test page.

[82] Then, in step 40, we proceed to the restoration of the content of the test page, that is to say the initial content which was backed up in the present case in the memory page with address 0 will be written back to the memory page that has just been the test page.

[83] It is then possible to change test page again and the cycle is repeated until all the pages of the memory have been used as test page. At that moment, as a function of the total number of accumulated errors, the refresh period is increased or decreased. This therefore marks the end of a main cycle of regulation of the refresh period. Then it is again the first page of the memory that becomes the test page and the operations described above are performed again in succession for all the pages of the memory so as to carry out a new cycle of regulation of the refresh period.

[84] As a variant, the modification of the refresh period can also be done as soon as a certain number of errors has been detected, without waiting for the end of the test of the entire memory if one wants to be able to react faster to a temperature variation for example.

[85] If one does not wish to place any constraint on the use of certain pages of the memory, it is possible to use a small external memory to temporarily back up the data of the

tested pages. Of course, this memory must then contain as many pages as are tested in an operating cycle of the algorithm.

[86] In the case where just a single test page is used, it is then particularly advantageous and particularly simple to use one of the buffer memories CH0 or CH1 for the temporary backing up of the data of the tested page.

[87] The variant of the invention illustrated diagrammatically in FIGURE 4 makes it possible to reduce the refresh frequency and hence the consumption in standby mode, even at maximum operating temperature. The principle of this variant comprises tagging, in the course of the first cycle of regulation of the refresh period of the memory, the pages having the least retention, of labeling them as "low" (step 45) and of subsequently refreshing them more often than the others (step 46). Specifically, generally, more than 90% of the pages of the memory can be refreshed 2 to 4 times less often than is necessary for the cells of minimum retention. Several modes of implementation of this variant are possible.

[88] More precisely, as far as the sequencing is concerned, it is conceivable to refresh all the pages at the maximum frequency, and to skip the refreshing of the pages that are not labeled as "low" one time out of two. Or else, the system can work at the minimum frequency plus 10% and insert an additional refreshing of the "low" page with address n just before (or after) processing the page with address (N/2+n) modulo N, where N denotes the total number of pages in the memory.

[89] As far as the storing of the so-called "low" page information is concerned, it is possible to use for example a static random access memory (SRAM memory) of N words. Each

word then makes it possible to label whether or not the page is low and possibly the state of its refresh if the first sequencing solution mentioned hereinabove is opted for.

[90] It is also possible to use an associative memory of P words with P equal to substantially 10% of N, this memory storing, for each page with address n labeled as "low," the value (N/2+n) modulo N. When the refresh address counter arrives at one of the addresses present in the associative memory, it then inserts a refresh of the page with address n.

[91] The choice of these two modes of storage depends on the relevant number of "low" pages. If this number is low, the second solution just described is more economical.

[92] As far as the decision algorithm is concerned, it is for example possible to refresh the memory at the frequency F and to test each page at the frequency F/2. When an error is detected, and if there are fewer than N "low" pages, the page for which an error is detected is labeled as "low" and will be refreshed at 2F and tested at F. If on the other hand there are already N low pages, this page is not labeled as low but the frequency F is increased. If, when all the pages have been tested, there have been no errors, F is decreased and the testing of the memory is begun again.

[93] Although preferred embodiments of the method and apparatus of the present invention have been illustrated in the accompanying Drawings and described in the foregoing Detailed Description, it will be understood that the invention is not limited to the embodiments disclosed, but is capable of numerous rearrangements, modifications and substitutions without departing from the spirit of the invention as set forth and defined by the following claims.